

15/3,K/1 (Item 1 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02713163 SUPPLIER NUMBER: 102205781 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Patent watch.

Microprocessor Report, 17, 4, 31(2)

April, 2003

ISSN: 0899-9341 LANGUAGE: English RECORD TYPE: Fulltext

WORD COUNT: 1195 LINE COUNT: 00117

... architecture for providing explicit multithreading.

6,463,529: Compaq-Processor based system with system wide reset and partial system reset capabilities.

6,463,580: Intel- **Parallel** processing utilizing highly correlated data values.

Issued: October 1, 2002

6,460,115: IBM-System and method for prefetching data to **multiple levels** of **cache** including selectively using a software hint to override a hardware prefetch mechanism.

6,460,116: AMD-Using separate **caches** for variable and generated fixed-length instructions.

6,460,129: Fujitsu-Pipeline operation method and pipeline operation device to interlock the translation of instructions based on the operation of a non-pipeline operation unit.

6,460,132: AMD-Massively **parallel** instruction predecoding.

6,460,134: Intrinsity-Method and apparatus for a late pipeline enhanced floating point unit.

Issued: September 24, 2002

6,457,121: Intel-Method and apparatus for reordering data in X86 ordering.

6,457,120: IBM-Processor and method including a cache having confirmation bits for improving address **predictable** branch instruction target predictions.

6,457,119: Intel-Processor instruction pipeline with error detection scheme.

6,457,118: Hitachi-Method and system for selecting and...

15/3,K/2 (Item 2 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02445591 SUPPLIER NUMBER: 66170678 (USE FORMAT 7 OR 9 FOR FULL TEXT)

PATENT WATCH. (News Briefs)

Belgard, Rich

Microprocessor Report, 14, 9, 42

Sept, 2000

ISSN: 0899-9341 LANGUAGE: English RECORD TYPE: Fulltext

WORD COUNT: 622 LINE COUNT: 00065

... data is available. During the same cycle that the reexecute signal is negated, result data is presented on the result bus.

6,067,616

Branch **prediction** device with **two levels** of branch prediction **cache**

Filed: April 26, 1996

Issued: May 23, 2000

Inventors: David Stiles et al.

Claims: 11

Assignee: AMD

An improved branch prediction **cache** scheme utilizes a hybrid **cache** structure. The **cache** provides **two levels** of branch information **caching**. A fully associative first level **cache** encaches full prediction information for a limited number of branch instructions. The second, direct-mapped **cache** encaches only partial prediction information, but does so for a much larger number of branch instructions. As each branch instruction is fetched and decoded, its address is used to perform parallel look-ups in the two branch prediction **caches**.

6,067,615

Reconfigurable processor for executing successive function sequences
in a processor operation

Filed: November 9, 1995

Issued: May 23, 2000

Inventor: Eric Upton...

15/3,K/3 (Item 3 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02423560 SUPPLIER NUMBER: 63676068 (USE FORMAT 7 OR 9 FOR FULL TEXT)

TI RELEASES DETAILS ON FIRST 'C64XX. (Product Announcement)

Eyre, Jennifer; Bier, Jeff

Microprocessor Report, 14, 6, 28

June, 2000

DOCUMENT TYPE: Product Announcement ISSN: 0899-9341 LANGUAGE:

English RECORD TYPE: Fulltext

WORD COUNT: 891 LINE COUNT: 00071

Because the 'C64xx is capable of very high levels of parallelism (for example, execution of four 16-bit multiplications in **parallel** with four 16-bit additions), getting data in and out of the processor core is a key concern--particularly given its high initial target clock...

...of much higher levels of parallelism and requires commensurately more data to operate at peak efficiency.

The 'C64xx chip will contain two level-one (L1) **caches**, one for data and one for instructions. Each of the L1 **caches** contains 16K of memory (four times that of the L1 **caches** on the 'C6211). The L1 program **cache** is direct mapped; the L1 data **cache** is two-way set-associative. The L1 **caches** are fed by a unified **level - two** (L2) **cache**, which contains four 32K memory banks totaling 128K (double the L2 of the 'C6211). The L2 **cache** can be configured as non-**cached** RAM, as a set-associative **cache**, or as a partitioned combination of the two. For example, the L2 can be configured as one bank of RAM, with the remaining three banks set up as a three-way set-associative **cache**. The L2 **cache** is fed via what TI calls an enhanced DMA controller, or EDMA. The EDMA supports 32 channels, and TI claims it can support more than...

15/3,K/4 (Item 4 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02187479 SUPPLIER NUMBER: 20813319 (USE FORMAT 7 OR 9 FOR FULL TEXT)

The State of Processors. (Intel's 80486 and Pentium series microprocessors, and Motorola 680x0 and PowerPC microprocessors) (PC Tech Tutor)

(Technology Information) (Column)

Barrett, Neil

Magazine, v17, n14, p307(3)

July, 1998

DOCUMENT TYPE: Column ISSN: 0888-8507 LANGUAGE: English

RECORD TYPE: Fulltext; Abstract

WORD COUNT: 3837 LINE COUNT: 00287

... the 80486's single pipeline, although for best performance, programs had to be optimized so that the pipelines would work together. Second, it employed branch- **prediction** technology to help minimize the delays often incurred when a branch instruction alters the flow of instruction execution. Third, the Pentium increased the speed of...

...MHz memory bus (initially a 60-MHz bus); the 486 used a 33-MHz version. Fourth, the Pentium came with built-in power management. Fifth, **two separate Level 1 caches**--one for data and the other for instructions--allowed programs to be optimized fully in both categories, and a separate floating-point pipeline improved the...

...two-pipeline technology. It improved pipelining by offering the results of instructions to both pipelines at once to reduce stalling, and by including better branch **prediction** and out-of-order completion (allowing instructions to be executed out of their coded order, as long as they don't rely on results from...Intel introduced the Pentium with MMX, a set of 57 additional instructions designed to improve the multimedia capabilities of the Pentium. These instructions focus on **parallel** execution and employ a technique called single instruction, multiple data (SIMD) to do their work. As the name suggests, SIMD allows a single instruction to work with more than one piece of data at the **same time**, thereby allowing the instruction to produce results more quickly. But this wasn't the only change in the MMX-adorned Pentiums. The pipeline increased to six stages from five, the **two Level 1 caches** were each increased from 8K to 16K, and branch **prediction** was improved.

The Pentium Pro and Pentium II

More than a year before the Pentium with MMX hit the market, Intel introduced its successor to...

15/3,K/5 (Item 5 from file: 275)
DIALOG(R)File 275:Gale Group Computer DB(TM)
(c) 2004 The Gale Group. All rts. reserv.

02110258 SUPPLIER NUMBER: 19840758 (USE FORMAT 7 OR 9 FOR FULL TEXT)
A RAM primer. (PC Tech: Tutor) (Technology Information) (Tutorial)
Randall, Neil
PC Magazine, v16, n18, p267(2)
Oct 21, 1997
DOCUMENT TYPE: Tutorial ISSN: 0888-8507 LANGUAGE: English
RECORD TYPE: Fulltext; Abstract
WORD COUNT: 2536 LINE COUNT: 00186

... 50 ns. The disadvantage is that SRAM is much more expensive than DRAM. SRAM's most common use in PCs is in the second-level **cache**, also called the L2 cache.

L2 cache . Caching is the art of predicting what data will be requested next and having that data already in hand, thus speeding execution. When your CPU makes a data request, the data can be found in one of four places: the L1 **cache**, the L2 **cache**, main memory, or in a physical storage system (such as a hard disk). L1 **cache** exists on the CPU, and is much smaller than the other three. The L2 **cache** (second-level **cache**) is a **separate** memory area, and is configured with SRAM. Main memory is much larger and consists of DRAM, and the physical storage system is much larger again but is also much, much slower than the other storage areas. The data search begins in the L1 **cache**, then moves out to the L2 **cache**, then to DRAM, and then to physical storage. Each level consists of progressively slower components. The function of the L2 cache is to stand between DRAM and the CPU, offering faster access than DRAM but requiring sophisticated **prediction** technology to make it useful. The term cache hit refers to a successful location of data in L2, not L1. The purpose of a cache...

...the days of the 386, and is still in place in the L2 cache of many PCs. It's called asynchronous because it's not **in sync** with the system clock, and therefore the CPU must wait for data requested from the L2 cache. The wait isn't as long as it...

15/3,K/6 (Item 6 from file: 275)
DIALOG(R)File 275:Gale Group Computer DB(TM)
(c) 2004 The Gale Group. All rts. reserv.

02104511 SUPPLIER NUMBER: 19758435 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Four-way superscalar PA-RISC processors. (HP PA 8000 and HP PA 8200)
(Product Information)
Scott, Anne P.; Burkhardt, Kevin P.; Kumar, Ashok; Blumberg, Richard M.; Ranson, Gregory L.
Hewlett-Packard Journal, v48, n4, p8(8)

August, 1997

ISSN: 0018-1153

LANGUAGE: English

RECORD TYPE: Fulltext; Abstract

WORD COUNT: 6695

LINE COUNT: 00526

... prediction convention should be followed or whether the opposite convention should be used. Compilers using either heuristic methods or profile-based optimization can use static **prediction** mode to communicate branch probabilities effectively to the hardware.

Cache Design

The PA 8000 features large, single-level, off-chip, direct-mapped instruction and data caches. Both caches support configurations of up to four megabytes using industry-standard **synchronous** SRAMs. Two complete copies of the data cache tags are provided so that two independent accesses can be accommodated and need not be to the same cache line.

Why did we design the processor without on-chip **caches** ? The main reason is performance. Competing designs incorporate small on-chip **caches** to enable higher clock frequencies. Small on-chip **caches** support benchmark performance but fade on large applications, so we felt we could make better use of the die area. The sophisticated IRB allows us to hide the effects of a pipelined two-state **cache** latency. In fact, our simulations demonstrated only a 5% performance improvement if the **cache** were on-chip and had a single-cycle latency. The flat **cache** hierarchy also eliminates the design complexity associated with a **two - level cache** design.

Chip Statistics

The PA 8000 is fabricated in HP's 0.5-micrometer, 3.3-volt CMOS process. Although the drawn geometries are not...

15/3,K/7 (Item 7 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02087522 SUPPLIER NUMBER: 19652651 (USE FORMAT 7 OR 9 FOR FULL TEXT)
AIM partnership unveils new chips. (Apple, IBM, Motorola) (IBM and Motorola debut PowerPC 750, new versions of Mach 5 processors) (Product Announcement)

Morgenstern, David

MacWEEK, v11, n30, p1(2)

August 4, 1997

DOCUMENT TYPE: Product Announcement ISSN: 0892-8118 LANGUAGE:
English RECORD TYPE: Fulltext; Abstract

WORD COUNT: 575 LINE COUNT: 00047

... end of October, is the first PowerPC processor tuned for performance with the Mac OS. The chip offers a range of technologies such as advanced **cache** architecture and power-saving modes, letting it run on high-performance desktops and notebooks.

Arthur's greatest difference from **past** PowerPC designs is a built-in **Level 2 cache** controller with a **separate** backside **cache** bus. The architecture lets the chip process instructions, send data over the system bus and communicate with the **cache** **simultaneously** .

In addition, **cache** tags (pointers to data on the **cache**) are stored within the 750 instead of on the **cache** itself; this lets the chip decide more quickly if it needs to fetch data from the **cache** and eliminates unnecessary trips to the **cache** .

The 750 is produced with a new 0.25-micron process, which lowers power requirements to about 5.7 watts, analysts said. It also integrates...

15/3,K/8 (Item 8 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

01704238 SUPPLIER NUMBER: 16255432 (USE FORMAT 7 OR 9 FOR FULL TEXT)
DEC's new Alpha MPU samples at 266/300MHz. (Alpha 21164 microprocessor)

Krause, Reinhardt

Electronic News (1991), v40, n2031, p2(2)

Sept 12, 1994

ISSN: 1061-6624

LANGUAGE: ENGLISH

RECORD TYPE: FULLTEXT; ABSTRACT

WORD COUNT: 936

LINE COUNT: 00072

...ABSTRACT: 5,000 units. The price of the microprocessor is considered high, but DEC officials say that they have introduced processors with high prices in the **past** and then lowered prices later in the product cycle. The 21164 has 9.3 million transistors and includes a **two - level**, on-chip **cache** memory. Cray Research is planning to use the 21164 in its next massively **parallel** processing supercomputer and two other companies, Aspen and Carrera, may also be using the processor. DEC is hoping that the company can use the 21164...

15/3,K/9 (Item 9 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

SUPPLIER NUMBER: 07361534 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Choose the right **parallel** architecture; **shared memories** are getting the nod for technical computing. (includes related article on distributed architectures)

Curran, Lawrence

Electronic Design, v37, n11, p41(6)

May 25, 1989

ISSN: 0013-4872

LANGUAGE: ENGLISH

RECORD TYPE: FULLTEXT; ABSTRACT

WORD COUNT: 3212

LINE COUNT: 00261

... Alliant, Convex, Sequent, Apollo Computer Inc., Ardent Computer Corp., Digital Equipment Corp., Pyramid Technology Corp., and Stellar Computer Inc.

Bartlett points out that shared-memory **parallel** multiprocessing almost always triggers a need for some kind of **cache** memory to ensure that a large percentage of processor references to memory go to **cache** storage instead of to main memory, conserving bus bandwidth. Encore's Unix-based Multimax 500 family of symmetrical multiprocessors, introduced this **past** February, uses the National Semiconductor NS32532 microprocessor and employs **two levels** of **cache**.

With **two** NS32532s per processor board, the Multimax 500 offers 17 MIPS per board--170 MIPS in the maximum system configuration of 10 boards. The first **cache** level is integrated into the NS32532 and consists of 1.5 kbytes of instruction and data **cache** to speed I/O operations. In addition, there are 256 kbytes of direct-mapped **cache** per CPU that employ a write-deferred protocol to reduce system bus drain. This second **cache** level is instrumental in preventing many memory references from falling through to the main memory. Bartlett says the writer-deferred protocol results in the CPU...

15/3,K/10 (Item 1 from file: 621)

DIALOG(R)File 621:Gale Group New Prod. Annou. (R)

(c) 2004 The Gale Group. All rts. reserv.

Supplier Number: 83443686 (USE FORMAT 7 FOR FULLTEXT)

Microsoft Delivers World-Record Performance With the Availability of Windows Datacenter Server Limited Edition and SQL Server 2000 Enterprise Edition.

PR Newswire, pSFM04104032002

March 4, 2002

Language: English Record Type: Fulltext

Document Type: Newswire; Trade

Word Count: 608

... 4.6C, a 32-way Unisys ES7000 server with 32 Intel Pentium III Xeon 900 MegaHertz processors, 12 gigabytes of memory and a 2 megabyte **level two cache** supported 26,000 **concurrent** SD benchmark users -- more than ever before -- with an average dialog response time of 1.97 seconds per request, a world record for the benchmark. This result equates to 2,606,000

fully processed business line items per hour, and also marks the second time in the **past** six months that Windows and SQL Server have been teamed to claim the top spot. The benchmark complies fully with the SAP Benchmark Council's...

15/3,K/11 (Item 1 from file: 636)
DIALOG(R)File 636:Gale Group Newsletter DB(TM)
(c) 2004 The Gale Group. All rts. reserv.

02701107 Supplier Number: 45478641 (USE FORMAT 7 FOR FULLTEXT)
STORAGE: **CAMBEX ENTERS MAINFRAME DASD MARKET WITH LOW-COST DISK ARRAY;**
COMPANY'S FIRST CASCADE ARRAY DASD OFFERS FULLY MIRRORED STORAGE SYSTEMS
FOR PRICE OF NON-MIRRORED PRODUCTS
EDGE: Work-Group Computing Report, v6, n256, pN/A
April 17, 1995
Language: English Record Type: Fulltext
Document Type: Newsletter; Trade
Word Count: 1757

... of these are: support for ESCON and parallel channels at all data rates, record or track level caching, dual copy data protection, remote dual copy, **concurrent** copy, transient dual copy to monitor the array, anticipate drive failure and automatically invoke a dual-copy spare volume, and DASD fast-write to enhance...

...operator training is fully protected and utilized.
To assist users in applying DASD, Cambex systems specialists help user configure Cascade arrays and Cambex 3990 controller **cache** memory to optimize storage subsystem performance for tasks from large batch jobs to real-time transaction processing. As the only independent manufacturer of large system **cache** products, Cambex provides this performance optimization at the lowest cost.

SUBSTANTIAL ERROR PROTECTION Cascade systems are designed for superior performance in IBM environments. **Parallel** processing arrangement of the disk drives, **multi - level caching**, **predictive** failure analysis and Cascade volume sparing and redundancy assure a level of continuous availability not possible with earlier generations of disk subsystems.

The systems also...

15/3,K/12 (Item 2 from file: 636)
DIALOG(R)File 636:Gale Group Newsletter DB(TM)
(c) 2004 The Gale Group. All rts. reserv.

01063112 Supplier Number: 40630791 (USE FORMAT 7 FOR FULLTEXT)
WORLD'S TECHNOLOGY PROVIDES DPS 9000 PERFORMANCE
Productivity Software, v2, n1, pN/A
Jan, 1989
Language: English Record Type: Fulltext
Document Type: Newsletter; Trade
Word Count: 1052

... in the SCU between the CPU and main memory is both large and fast, with 512KB capacity and 6.0 nanosecond access time. The system **cache** is organized with a four-level set associative format, and provides performance benefits by anticipating and buffering memory requirements for both the CPUs and the I/O processors. Both the system **cache** and the CPU **cache** support the pipeline architecture, providing high-speed access to operands and instructions.

The CPU's main memory has a capacity of up to one gigabyte, using one megabit technology, and has its own **two - level cache** of 128KB in each processor with access time of 3.0 nanoseconds. There is a 512KB processor **cache** maximum in a four processor system and up to one megabyte of system **cache** in a fully configured system.

Pipeline Processing

The DPS 9000 CPU incorporates a high performance seven-stage instruction pipeline. Up to seven different instructions can be executed

simultaneously in a single processor, and the pipeline design enables the CPU to begin executing instructions before the previous ones have been completed.

The DPS 9000 also uses a new branch **prediction** system which keeps instructions flowing smoothly through the pipeline, even across many transfer/branch instructions in program coding.

Vector Processing

Like the previously-announced Honeywell...

15/3,K/13 (Item 1 from file: 16)
DIALOG(R)File 16:Gale Group PROMT(R)
(c) 2004 The Gale Group. All rts. reserv.

07549276 Supplier Number: 63135463 (USE FORMAT 7 FOR FULLTEXT)
Supercomputers Make a Comeback; Proving that rumors of their demise were greatly exaggerated, very-high-performance computers show promise as tools to do the heavy lifting for e-commerce applications in mainstream IT shops. (Product Information)

Anthes, Gary H.

Computerworld, p54(1)

July 3, 2000

Language: English Record Type: Fulltext Abstract

Document Type: Tabloid; Trade

Word Count: 1383

... which rears its ugly head when processors wait idly for data from memory or, worse, disk.

A solution to the latency problem is to add **multiple levels** of **cache** storage on or near the processor chip where commonly used data or instructions can be retrieved very rapidly. Systems today have three levels of **cache**, but more will be added, Dongarra says.

"We have failed to capitalize on the performance potential of scalable, **parallel** machines," says Ken Kennedy, director of the Center for High Performance Software at Rice University in Houston. Programmers haven't been good enough at structuring their code for **parallel** processing and have had difficulty optimizing their code for the complex memory hierarchies in many **parallel** systems, he says.

Advances to Come

But Kennedy says research shows promise for shifting those burdens from programmers to compilers and other tools. Compilers will...

...memories and do more global optimization by considering entire programs rather than individual routines, he says. And higher bandwidth inside machines will reduce memory latency, he **predicts**.

Tera Computer Co. (now Cray Inc., having bought the Cray supercomputer business from Silicon Graphics Inc. in April) devised another solution to the latency problem...

15/3,K/14 (Item 2 from file: 16)
DIALOG(R)File 16:Gale Group PROMT(R)
(c) 2004 The Gale Group. All rts. reserv.

05961353 Supplier Number: 53234766 (USE FORMAT 7 FOR FULLTEXT)
Notebooks worth the weight. (IBM ThinkPad, HP OmniBook, Compaq Armada, Toshiba Tecra, Gateway Solo notebooks) (Hardware Review) (Evaluation)

Kvitka, Andre

InfoWorld, pNA(1)

Nov 16, 1998

Language: English Record Type: Fulltext

Article Type: Evaluation

Document Type: Magazine/Journal; Trade

Word Count: 2183

... the core, resulting in more power in less space. The chip also has the Dual Independent Bus (DIB) architecture, which combines a high-speed 512KB **Level 2 cache** with **two** independent system buses for

simultaneous parallel access to data and dynamic execution. As with the desktop version, the mobile 300-MHz Pentium II processor combines three data-processing techniques to manipulate data more intelligently and efficiently. These techniques **predict** and analyze software instructions to optimize processor workload. New to this version of the chip is a technology called Quick Start that drops power consumption...

15/3,K/15 (Item 3 from file: 16)
DIALOG(R)File 16:Gale Group PROMT(R)
(c) 2004 The Gale Group. All rts. reserv.

05676680 Supplier Number: 50157537 (USE FORMAT 7 FOR FULLTEXT)

IBM's System/390 servers set a new speed record

LIFEN, FLORENCE

Government Computer News, v17, n19, p44

June 29, 1998

Language: English Record Type: Fulltext

Article Type: Article

Document Type: Magazine/Journal; Tabloid; Trade

Word Count: 484

... vice president of S/390 global hardware development.

The module holds a 0.25-micron CMOS 6X central processor containing 25 million transistors, 8M of **Level 2 cache** containing 60 million transistors and **two** cryptographic coprocessors.

IBM has rated the 6X Turbo uniprocessor at 125 MIPS, twice the rating of the previous-generation CMOS uniprocessor.

IBM engineers have designed the S/390 G5 server with four system buses between the Level 2 **cache** and memory to handle mixed workloads, especially for Unix, Mauri said.

"Unix workloads and anything written in Java, C or C++ will perform significantly better on a S/390," Mauri said. IBM has licensed Lotus Domino Server running on the OS/390 to 160 customers in the **past** six months, officials said.

IBM also added 121 new instructions and hardware features to the S/390 G5 server to run standard binary floating-point calculations.

A new Geoplex clustering facility for the Fast Ethernet S/390 G5 servers extends **Parallel** Sysplex clustering over distances up to 45 kilometers, or about 25 miles. Mauri said IBM users asked for that much distance to satisfy their disaster...

15/3,K/16 (Item 4 from file: 16)
DIALOG(R)File 16:Gale Group PROMT(R)
(c) 2004 The Gale Group. All rts. reserv.

04058735 Supplier Number: 45906316 (USE FORMAT 7 FOR FULLTEXT)

The World According to INTEL

VARbusiness, p62

Nov 1, 1995

Language: English Record Type: Fulltext

Document Type: Magazine/Journal; Trade

Word Count: 2890

... in a single package," according to Intel. Taking a systems-level approach to chip layout, Intel has coupled the Pentium Pro die with a companion **level - two cache** die. The design essentially combines processor speed, high-speed memory and a blazingly fast bus in a single package.

"Think of the 486 as a...

...Pro is superscalar (like the original Pentium) and superpipelined, and it supports register renaming. It also includes support for Dynamic Execution, Intel's name for **simultaneous** multiple-branch **prediction**, data flow analysis and speculative execution. The Pentium Pro can crunch bigger chunks of data, allocate computational resources more efficiently and optimize work done in...

15/3,K/17 (Item 5 from file: 16)
DIALOG(R)File 16:Gale Group PROMT(R)
(c) 2004 The Gale Group. All rts. reserv.

03754743 Supplier Number: 45331534 (USE FORMAT 7 FOR FULLTEXT)
ISSCC moves beyond P6
Electronic Engineering Times, p1
Feb 13, 1995
Language: English Record Type: Fulltext
Document Type: Magazine/Journal; Trade
Word Count: 1376

... it is allowed to alter the 'official' register set and the caches.
To keep 30 instructions in the air with existing X86 programs requires looking **past** a lot of branch instructions. Branches, calls and so forth can occur as often as every four or five instructions in some X86 programs. To cope, the P6 has possibly the most advanced branch- **prediction** and speculative -execution facilities yet attempted, according to some sources. The chip can speculatively execute **past** five to seven branches, keeping track of its guesses and resolving them as instructions are retired.

Because the P6 will serve a market in which board designers vary in skill and technology, simply bringing a 133-MHz secondary **cache** bus to the pins of the PGA and leaving the details to the customer was not an option. Intel opted for a more controversial approach...
The P6 die (the CPU runs at 2.9 V) mounted on a ceramic substrate beside the CPU and wire-bonded to it. The separate **cache** die includes 256 kbytes of data SRAM and **cache** tags, operating under the control of a **level - two cache** controller on the CPU die.

Thus, the choice of **cache** SRAM technology, design of the controller, and implementation of the **cache** bus have all been decided by the time the board vendor receives the P6 PGA.

That approach meets the immediate needs of the P6, but it's clear from other ISSCC papers that it won't silence the debate about **cache** implementation.

Breakthrough research is being reported in two areas that will change the way secondary **cache** is handled. The first is the appearance of extremely fast **synchronous** SRAMs using wave pipelining. Instead of signals moving through the cache from latch to latch, remaining stable between clock transitions, data is propagated as pulses...

15/3,K/18 (Item 6 from file: 16)
DIALOG(R)File 16:Gale Group PROMT(R)
(c) 2004 The Gale Group. All rts. reserv.

03054777 Supplier Number: 44156572 (USE FORMAT 7 FOR FULLTEXT)
National taps V3 Corp. for DRAM controllers
Electronic Engineering Times, p3
Feb 11, 1993
Language: English Record Type: Fulltext
Document Type: Magazine/Journal; Trade
Word Count: 782

... business at Ramtron. 'Basically, we've put an SRAM cache in each DRAM chip. But there is a unique twist - we load the cache in **parallel** from the memory array.'

As a result, on a **cache** read miss, an entire 2-kbit page is loaded into the local SRAM in a single 35-ns cycle. So, further reads to that page ...

...precharge and refresh delays are effectively hidden.

'But the part does require some changes to the memory controller,' Oseth noted. 'If you already have a **level - two cache** controller, we need to add about 300 gates to fully exploit the EDRAM. Because of the difference, we've had to be very proactive with...

...in embedded designs, where locality of reference tends to be high, the

EDRAM could deliver performance equal or superior to that of systems with large **level - two caches**. Thus, V3's choice to target embedded systems could be fortuitous for Ramtron as well.

Lack of credibility

Both V3 and Ramtron have suffered in the **past** from lack of a credible second-source. V3's problem has been essentially solved by National, but Ramtron's remains. Oseth claimed, however, that the...

15/3,K/19 (Item 1 from file: 148)

DIALOG(R)File 148:Gale Group Trade & Industry DB
(c)2004 The Gale Group. All rts. reserv.

09646072 SUPPLIER NUMBER: 17776886 (USE FORMAT 7 OR 9 FOR FULL TEXT)
EDN's 22nd annual microprocessor/microcontroller directory. (Directory)

Levy, Markus; Leonard, James P.

ENR, v40, n19, p33(52)

Feb 14, 1995

DOCUMENT TYPE: Directory ISSN: 0012-7515 LANGUAGE: English

RECORD TYPE: Fulltext; Abstract

WORD COUNT: 18643 LINE COUNT: 01483

... million-transistor 21164 has a seven-stage integer pipeline that contains two integer units. It also includes a nine-stage floating-point unit that can **simultaneously** issue an add and a multiply in one cycle. A unique feature of the 21164 is the 96-kbyte on-chip, **level - two cache** that provides 50% better performance than an external **cache**. To help increase loading efficiency, the 21164 contains merge logic that looks ahead to see if more than one reference is being made to the same **cache** block. The merge logic can compress up to 20 loads in a row.

The 21064 has a seven-stage integer pipeline and a 10-stage floating-point pipeline and can issue up to two instructions per cycle. It implements branch-**prediction** hardware based on a single branch **history** bit (2 bits in the 21064A) stored with each instruction in the instruction cache. The chip's pipelines are static and dynamic. The first four...and decode. Three additional stages have been added to the integer pipe to make it symmetrical with the floating-point pipe. This architecture simplifies pipeline **synchronization** and exception handling; it also eliminates the need for a floating-point queue. The CPU's pipeline encompasses two integer ALUs, five floating-point graphic units, and a load/store unit. Sun also includes a 2-bit dynamic branch-**prediction** mechanism, which is part of its prefetch unit. As the 16-kbyte instruction **cache** fills, the CPU uses two extra bits per instruction to tag on information related to the branch **prediction** for that instruction.

UltraSPARC uses data buffers to isolate the **level - two cache** from the system bus. These buffers enable overlapping of system transactions and perform error detection/correction. The processor contains an on-chip **level - two cache** controller, and the system bus can run at one-half to one-third the processor frequency. Sun claims that instructions and data can pass between the **level - two cache** and the CPU at 2.6 Gbytes/sec.

Special instructions: SPARC V9 adds several instructions to the V8 specification: conditional move, 64-bit integer multiply...

15/3,K/20 (Item 2 from file: 148)

DIALOG(R)File 148:Gale Group Trade & Industry DB
(c)2004 The Gale Group. All rts. reserv.

09644235 SUPPLIER NUMBER: 16821865 (USE FORMAT 7 OR 9 FOR FULL TEXT)
ISSC moves beyond P6. (International Solid-State Circuits

Conference) (design for Intel's P6 processor)

Wilson, Ron

Electronic Engineering Times, n835, p1(2)

Feb 13, 1995

ISSN: 0192-1541 LANGUAGE: English RECORD TYPE: Fulltext; Abstract

WORD COUNT: 1480 LINE COUNT: 00117

... allowed to alter the "official" register set and the caches.

Speculative execution

To keep 30 instructions in the air with existing X86 programs requires looking **past** a lot of branch instructions. Branches, calls and so forth can occur as often as every four or five instructions in some X86 programs. To cope, the P6 has possibly the most advanced branch- **prediction** and speculative-execution facilities yet attempted, according to some sources. The chip can speculatively execute **past** five to seven branches, keeping track of its guesses and resolving them as instructions are retired.

Because the P6 will serve a market in which board designers vary in skill and technology, simply bringing a 133-MHz secondary **cache** bus to the pins of the PGA and leaving the details to the customer was not an option. Intel opted for a more controversial approach...

...3-V die (the CPU runs at 2.9 V) mounted on a ceramic substrate beside the CPU and wire-bonded to it. The separate **cache** die includes 256 kbytes of data SRAM and **cache** tags, operating under the control of a **level - two cache** controller on the CPU die.

Thus, the choice of **cache** SRAM technology, design of the controller, and implementation of the **cache** bus have all been decided by the time the board vendor receives the P6 PGA.

That approach meets the immediate needs of the P6, but it's clear from other ISSCC papers that it won't silence the debate about **cache** implementation.

Breakthrough research is being reported in two areas that will change the way secondary **cache** is handled. The first is the appearance of extremely fast **synchronous** SRAMs using wave pipelining. Instead of signals moving through the cache from latch to latch, remaining stable between clock transitions, data is propagated as pulses...

15/3,K/21 (Item 3 from file: 148)

DIALOG(R)File 148:Gale Group Trade & Industry DB
(c)2004 The Gale Group. All rts. reserv.

08333459 SUPPLIER NUMBER: 17746378 (USE FORMAT 7 OR 9 FOR FULL TEXT)

The world according to Intel. (includes interview with Intel CEO Andrew Grove and related articles on Intel's new channel programs and VAR Triad Systems Corp) (Company Business and Marketing) (Cover Story)

Doyle, T. C.

VARbusiness, v11, n17, p62(8)

Nov 1, 1995

DOCUMENT TYPE: Cover Story ISSN: 0894-5802 LANGUAGE: English

RECORD TYPE: Fulltext; Abstract

WORD COUNT: 4682 LINE COUNT: 00367

... in a single package," according to Intel. Taking a systems-level approach to chip layout, Intel has coupled the Pentium Pro die with a companion **level - two cache** die. The design essentially combines processor speed, high-speed memory and a blazingly fast bus in a single package.

"Think of the 486 as a...

...Pro is superscalar (like the original Pentium) and superpipelined, and it supports register renaming. It also includes support for Dynamic Execution, Intel's name for **simultaneous** multiple-branch **prediction**, and flow analysis and speculative execution. The Pentium Pro can crunch bigger chunks of data, allocate computational resources more efficiently and optimize work done in...

15/3,K/22 (Item 4 from file: 148)

DIALOG(R)File 148:Gale Group Trade & Industry DB
(c)2004 The Gale Group. All rts. reserv.

17/99074 SUPPLIER NUMBER: 16802250 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Cambex enters mainframe DASD market with low-cost disk array; Company's first Cascade array DASD offers fully mirrored storage systems for price

of non-mirrored products.

Business Wire, p4110204

April 11, 1995

LANGUAGE: ENGLISH RECORD TYPE: FULLTEXT

WORD COUNT: 1831 LINE COUNT: 00154

... of these are: support for ESCON and parallel channels at all data rates, record or track level caching, dual copy data protection, remote dual copy, **concurrent** copy, transient dual copy to monitor the array, anticipate drive failure and automatically invoke a dual-copy spare volume, and DASD fast-write to enhance...

...operator training is fully protected and utilized.

To assist users in applying DASD, Cambex systems specialists help user configure Cascade arrays and Cambex 3990 controller **cache** memory to optimize storage subsystem performance for tasks from large batch jobs to real-time transaction processing. As the only independent manufacturer of large system **cache** products, Cambex provides this performance optimization at the lowest cost.

Substantial Error Protection

Cascade systems are designed for superior performance in IBM environments. **Parallel** processing arrangement of the disk drives, **multi-level caching**, **predictive** failure analysis and Cascade volume sparing and redundancy assure a level of continuous availability not possible with earlier generations of disk subsystems.

The systems also...

15/3,K/23 (Item 5 from file: 148)

DIALOG(R)File 148:Gale Group Trade & Industry DB

(c)2004 The Gale Group. All rts. reserv.

07571989 SUPPLIER NUMBER: 16423262 (USE FORMAT 7 OR 9 FOR FULL TEXT)

Superscalar RISC micronPs scale performance heights.

Levy, Markus

EDN, v39, n22, p18(2)

Oct 27, 1994

ISSN: 0012-7515 LANGUAGE: ENGLISH RECORD TYPE: FULLTEXT; ABSTRACT

WORD COUNT: 1339 LINE COUNT: 00104

... complete, but, until graduation, its results are tentative and may be discarded.

Unlike the 21164, the T5 has large (32-kbyte) primary-instruction and data **caches**. The processor supports 512 kbytes to 16 Mbytes of external secondary cache and uses a 128-bit data path to move data. The level-two...

...the system bus can run at one-half or one-third of the processor frequency. Sun states that instructions and data can pass between the level - two **cache** and the CPU at 2.6 Gbytes/sec.

UltraSPARC is estimated to run at 200 to 300 SPECint92 and 250 to 350 SPECfp92 at 167...in-order graduation.

Unlike other PowerPC CPUs, the 620 can perform both static and dynamic and branch prediction. Another difference is the CPU's huge **caches**. The 601 has a 32-kbyte unified **cache**, and the 620 implements a Harvard architecture with 32-kbyte instruction and data **caches**. Both **caches** are eight-way set-associative, and the data **cache** handles write-through and write-back modes.

Another key feature of the 620 is the 128-bit level - two **cache** interface. The unified level - two **cache** can be clocked at one, one-half, or one-third the frequency of the processor clock. The CPU also has a 40-bit address bus...

15/3,K/24 (Item 1 from file: 647)

DIALOG(R)File 647:CMP Computer Fulltext

(c) 2004 CMP Media, LLC. All rts. reserv.

01076500 CMP ACCESSION NUMBER: EBN19951211S0029

**PNY Sets Up SV Facility In An Effort To Penetrate Custom Module Market -
Multifaceted Plan Focuses On High Value-Added End**

Hugh G. Willett
ELECTRONIC BUYER'S NEWS, 1995, n 985, PG12
PUBLICATION DATE: 951211
JOURNAL CODE: EBN LANGUAGE: English
RECORD TYPE: Fulltext
SECTION HEADING: Semiconductor
WORD COUNT: 455

... We see a big opportunity in OEM product sales where we design and manufacture standard products for OEMs as well as custom derivatives such as **level - two cache** modules for PCs and flash cards for networking applications," Portnoy said.

The explosion in new memory technologies such as flash, burst extended data out, and **synchronous** DRAMs will make the memory sourcing and design task more complex, creating more opportunity for PNY, he said.

"**Synchronous** memories are new and there isn't a lot of expertise available yet," he said. "OEMs and semiconductor companies are coming to us and we have several **synchronous** projects going right now."

PNY will have to add additional engineering talent to exploit the opportunity, he said. Moving from a manufacturing mind-set to...

...is going to be a challenge, according to Portnoy.

Many of the module makers that have grown fat on the contract manufacturing profits of the **past** few years may fall out of the market if they mismanage their move up the value chain, he added.

15/3,K/25 (Item 2 from file: 647)
DIALOG(R)File 647:CMP Computer Fulltext
(c) 2004 CMP Media, LLC. All rts. reserv.

01069741 CMP ACCESSION NUMBER: VAR19951101S0038
The World According to INTEL (TAKING CARE OF BUSINESS)
T.C. Doyle
VARBUSINESS, 1995, n 1117, PG62
PUBLICATION DATE: 951101
JOURNAL CODE: VAR LANGUAGE: English
RECORD TYPE: Fulltext
SECTION HEADING: Cover Story
WORD COUNT: 2884

... in a single package," according to Intel. Taking a systems-level approach to chip layout, Intel has coupled the Pentium Pro die with a companion **level - two cache** die. The design essentially combines processor speed, high-speed memory and a blazingly fast bus in a single package.

"Think of the 486 as a...

...Pro is superscalar (like the original Pentium) and superpipelined, and it supports register renaming. It also includes support for Dynamic Execution, Intel's name for **simultaneous** multiple-branch **prediction**, data flow analysis and speculative execution. The Pentium Pro can crunch bigger chunks of data, allocate computational resources more efficiently and optimize work done in...

15/3,K/26 (Item 3 from file: 647)
DIALOG(R)File 647:CMP Computer Fulltext
(c) 2004 CMP Media, LLC. All rts. reserv.

01043865 CMP ACCESSION NUMBER: EBN19950220S0004
Intel Spins P6 Terms
Jonathan Cassell
ELECTRONIC BUYER'S NEWS, 1995, n 943, PG1
PUBLICATION DATE: 950220
JOURNAL CODE: EBN LANGUAGE: English

RECORD TYPE: Fulltext
SECTION HEADING: News
WORD COUNT: 1135

... typical)
The P6 is packaged in a multichip module (MCM) that contains the 3.5-million transistor microprocessor die as well as a 256-kilobyte **level - two cache** that can run at the full speed of the microprocessor. The **cache** is four-way set-associative, consisting of an Intel- manufactured **synchronous** burst static RAM.

Intel said it will begin volume production of the P6 in the second half of 1995. The company declined to reveal pricing or precise manufacturing volume in 1995.

Key to the P6's performance leap over the Pentium is its use of multiple branch **prediction**, data flow analysis, and speculative execution, the features which Intel feels justify the term dynamic execution. However, such techniques have already been applied in RISC...

15/3,K/27 (Item 4 from file: 647)
DIALOG(R)File 647:CMP Computer Fulltext
(c) 2004 CMP Media, LLC. All rts. reserv.

EBN 925 CMP ACCESSION NUMBER: EBN19941010S0072
Rattles Modules - As Demand For Second-Leve; Cache Modules Takes Off, Tight Supplies Of Fast SRAMs May Pose Glitch (Briefs)
Bettyann Tinnelly
ELECTRONIC BUYER'S NEWS, 1994, n 925, PG30
PUBLICATION DATE: 941010
JOURNAL CODE: EBN LANGUAGE: English
RECORD TYPE: Fulltext
SECTION HEADING: Product Focus
WORD COUNT: 1472

... RAMs can cost two to three times as much as asynchronous RAMs, and they can be difficult to get, Herbst said.

MMTG has been shipping **synchronous** burst cache modules since the fourth quarter of 1993, said Bob Bolger, fast SRAM module manager for Motorola. ``We've had to do the fastest ramp-up in our **history** since introducing our two latest products in the burst SRAM family: 32-K18, and 64-K18.''

The company has recently ramped up capacity for fast SRAMs, including **synchronous** burst devices, at two fabrication facilities in Austin and one in Scotland.

Cypress plans to introduce **synchronous** burst cache for PCs in the fourth quarter. **Synchronous** burst **cache** modules are generally considered a specialty product for workstations.

MicroModule Systems Inc. is taking a different approach to the same end by making and selling multichip modules for **cache**. MCMs place multiple unpackaged IC dice into a package the size of a traditional monolithic. ``Although our MCMs appear to be monolithic, they really consist...

...vice president of sales for MicroModule, Cupertino, Calif. ``It's a very high-density package-a 28 -mm square 160-pin QFP.''

The company provides **level - two cache** MCMs with pipeline burst SRAMs performing at speeds of 9 ns. MicroModule targets high-end workstation and file server customers such as Compaq, NCR, AST, IBM, and Dell. MicroModule has relationships with both Motorola and Micron for supplies of the 64-K18 dice it uses to construct its **cache** MCMs, Best said. As a result of those alliances, the company has not experienced the same shortages felt by traditional module makers. ``Lead time for...

15/3,K/28 (Item 1 from file: 674)
DIALOG(R)File 674:Computer News Fulltext
(c) 2004 IDG Communications. All rts. reserv.

085605

Supercomputers Make a Comeback

Proving that rumors of their demise were greatly exaggerated, very-high-performance computers show promise as tools to do the heavy lifting for e-commerce applications in mainstream IT shops. By Gary H. Anthes

Byline: Gary H. Anthes

Journal: Computerworld

Page Number: 54

Publication Date: July 03, 2000

Word Count: 1377

Line Count: 131

Text:

... which rears its ugly head when processors wait idly for data from memory or, worse, disk.

A solution to the latency problem is to add **multiple levels** of **cache** storage on or near the processor chip where commonly used data or instructions can be retrieved very rapidly. Systems today have three levels of **cache**, but more will be added, Dongarra says.

"We have failed to capitalize on the performance potential of scalable, **parallel** machines," says Ken Kennedy, director of the Center for High Performance Software at Rice University in Houston. Programmers haven't been good enough at structuring their code for **parallel** processing and have had difficulty optimizing their code for the complex memory hierarchies in many **parallel** systems, he says.

Advances to Come

But Kennedy says research shows promise for shifting those burdens from programmers to compilers and other tools. Compilers will...

... memories and do more global optimization by considering entire programs rather than individual routines, he says. And higher bandwidth inside machines will reduce memory latency, he **predicts**.

Tera Computer Co. (now Cray Inc., having bought the Cray supercomputer business from Silicon Graphics Inc. in April) devised another solution to the latency problem...

?